Virtual Adversarial Training: a Regularization Method for Supervised and Semi-supervised Learning

Takeru Miyato*,†,‡, Shin-ichi Maeda†, Masanori Koyama†,§ and Shin Ishii†,‡

arXiv:1704.03976v1 [stat.ML] 13 Apr 2017

Abstract—We propose a new regularization method based on virtual adversarial loss: a new measure of local smoothness of the output distribution. Virtual adversarial loss is defined as the robustness of the model's posterior distribution against local perturbation around each input data point. Our method is similar to adversarial training, but differs from adversarial training in that it determines the adversarial direction based only on the output distribution and that it is applicable to a semi-supervised setting. Because the directions in which we smooth the model are *virtually* adversarial, we call our method virtual adversarial training (VAT). The computational cost of VAT is relatively low. For neural networks, the approximated gradient of virtual adversarial loss can be computed with no more than two pairs of forward and backpropagations. In our experiments, we applied VAT to supervised and semi-supervised learning on multiple benchmark datasets. With additional improvement based on entropy minimization principle, our VAT achieves the state-of-the-art performance on SVHN and CIFAR-10 for semi-supervised learning tasks.

1 Introduction

In practical regression and classification problems, the sample size that can be used to tune the parameters of a predictive model is finite, and any training process is certain to suffer some overfitting; the training error will always be different from the test error [2], [42]. Regularization is a concept used to manage this inevitable gap between the training error and test error. It is often achieved by augmenting the loss function with a so-called regularization term, which prevents the model from overfitting to the loss function evaluated at finite sample points.

A Bayesian interpretation of the regularization term is a prior distribution that reflects our *a priori* knowledge or belief regarding the model. A popular *a priori* belief based on widely observed facts is that the outputs of most naturally occurring systems are smooth with respect to spatial and/or temporal inputs.

In fact, *smoothing the model* often works to our advantage in practice. For example, label propagation [44] is an algorithm that speculates on the class of unlabeled training samples based on the belief that close input data points tend to have similar class labels. Additionally, for neural networks (NNs), it is known that generalization performance can be improved by training the model to produce similar set of outputs to the set of inputs generated by adding random perturbation to each training input data point [4]. Several studies have also confirmed that this philosophy of *making the predictor robust against random and local perturbation* is also effective in semi-supervised learning [3], [21], [31], [41].

- *Contact :* takeru.miyato@gmail.com
- * *Preferred Networks, Inc., Tokyo, Japan,* † *Graduate school of Informatics, Kyoto University, Kyoto, Japan,* ‡ *ATR Cognitive Mechanisms Laboratories, Kyoto, Japan,* § *Department of Mathematical Science, Ritsumeikan University, Kyoto, Japan*
- *Shin-ichi Maeda currently belongs to Preferred Networks, Inc.*

However, the authors of [11], [37] found a weakness in this philosophy. They found that standard isotropic smoothing via random noise and random data augmentation often leaves the predictor particularly vulnerable to a small perturbation in a specific direction, that is, the adversarial direction, which is the direction in the input space to which the label probability of the model is most sensitive. The authors experimentally verified that the predictors trained with the standard isotropic regularization technique are likely to suffer from mistake on signals perturbed in the adversarial direction, even when the norm of the perturbation is so small that it cannot be perceived by human eyes.

Inspired by this finding, Goodfellow et al. [11] developed *adversarial training* that trains the model to assign similar labels to the training input data point as the neighborhood inputs in the adversarial direction. This attempt succeeded in improving generalization performance and made the model robust against adversarial perturbation. Goodfellow et al.'s work suggests that the locally isotropic model cannot be achieved via isotropic smoothing. Instead, one must strive to promote the local isotropy by investigating local anisotropy around each input data point and smoothing the model in the most anisotropic direction.

Our proposed regularization technique is a method that trains the output of the model to be isotropically smooth around each input data point irrespective of the availability of the training label. Thus, our algorithm is a natural extension of the past efforts to make the model isotropically smooth locally around each input data point. As inferred by [37], it is difficult to achieve local isotropic smoothness by applying input independent regularization on the parameters. This is especially true for the models with complex nonlinear relations between inputs and outputs. The aim of our approach is to improve generalization performance by resolving the complex anisotropic sensitivity of such models around each data point.

We define local distributional smoothness (LDS) at each input data point as information theoretic robustness of the model against the perturbation that can most greatly alter the output distribution. In this paper, we refer to this most dangerous perturbation as *virtual adversarial perturbation*. We propose a novel training method that maximizes the likelihood of the model, while promoting the model's LDS on each training input data point. We call this method *virtual adversarial training* (VAT).

VAT is a method that replaces the true label y in the original adversarial training with the posterior probability $p(y|x, \hat{\theta})$. The following list summarizes the advantages of this new method:

- applicability to semi-supervised learning tasks
- applicability to parametric models on which we can evaluate the gradient with respect to input and parameter
- small number of hyperparameters
- parametrization invariant regularization

The second advantage is worth emphasizing. At first glance, our algorithm appears as if it needs to solve an internal optimization problem to determine the virtual adversarial direction. For models such as NNs for which we can evaluate the gradient of the output with respect to the input, however, virtual adversarial perturbation admits an approximation that can be computed efficiently with the power method. This property enables us to implement VAT for NNs with no more than three times the computational cost of standard, regularization-free training.

Parametrization invariance is an important aspect of our algorithm because it is the property that makes the objective of our regularization unambiguous. The effects of many popular regularization methods, such as L_q regularization, can change with parametrization, and the analysis of this change can be difficult because of the models' nonlinear dependence on the parameters. This can be a challenge for developing and evaluating regularization methods. By contrast, our algorithm is parametrization invariant and there is no ambiguity in its effect; that is, as a function with respect to the input, the model being considered *optimal* by our regularized objective function does not change with reparametrization.

When we applied VAT to the supervised and semi-supervised learning of the permutation invariant task for the MNIST dataset, our method outperformed all contemporary methods other than some cutting-edge methods that use generative models. We also applied our method to semi-supervised learning on CIFAR-10 and Street View House House Numbers (SVHN) datasets, and confirmed that our method achieved superior or comparable performance to that of state-of-the-art methods.

This article also extends our earlier work [27] in four aspects:

- clarification of the objective function
- comparison between VAT and random perturbation training (RPT)
- additional set of extensive experiments
- evaluation of the virtual adversarial examples
- enhancement of the algorithm with entropy minimization

2 Related Works

Many classic regularization methods for NNs introduce random perturbation to input or hidden layers [4], [9], [35]. An early work by Bishop [4] showed that Gaussian perturbed input during training is equivalent to adding an extra regularization term to an original objective function derived from clean input data points. For small perturbations, the induced regularization term behaves similarly to a class of Tikhonov regularizers [39]. The introduction of random perturbation to inputs can be interpreted as a regularization aimed at smoothing the input-output relation of the NN. Another way to smooth the input-output relation is to impose constraints on the derivatives. For example, constraints may be imposed on the Frobenius norm of the Jacobian of the output with respect to the input. This approach was taken by Gu and Rigazio [13] in their deep contractive network. Instead of computing the computationally expensive full Jacobian, however, they approximated the Jacobian by the sum of the Frobenius norms of the layer-wise Jacobians computed for all adjacent pairs of hidden layers. The reduction in the test error achieved by the deep contrastive network, however, was not particularly significant. This is possibly because of their layer-wise approximation.

Dropout [35] is another popular method for regularizing NNs with random noise. Dropout is a method that randomly masks input or hidden variables in the network during training. From a Bayesian perspective, dropout is a method to introduce prior distribution for the parameters rather than a regularization that introduces perturbation to the inputs and the hidden variables [6], [26]. In this interpretation, dropout is a method that regularizes the model via Bayesian model ensembles, and is complementary to our approach, which directly augments the function with a regularization term.

Adversarial training was originally proposed by [37]. They showed that several NNs, including state-of-the-art NN models, are vulnerable to a very small perturbation in the direction to which the models' label assignment y of x is sensitive (adversarial), even when the perturbation is so small that human eyes cannot discern the difference. They also showed that training the models to be robust against adversarial perturbation is effective in reducing the test error. However, the definition of adversarial perturbation in [37] required a computationally expensive inner loop to identify the adversarial direction. To overcome this problem, Goodfellow et al. [11] proposed another definition of adversarial perturbation that admits a form of approximation that is free of the expensive inner loop (see the next section for details).

Bachman et al. [3] studied random perturbation in the setting of semi-supervised learning. The aim of their method was to minimize the prediction variance derived from the noise in the input layer. Our method can be seen as a variant of their method in which the isotropic, Gaussian random perturbation is replaced by the perturbation to which the model is most sensitive. As we show later, in our experiments, this replacement worked to the advantage of the model's generalization performance. Random image augmentation is a variant of random perturbation that simply augments the dataset with images perturbed by regular

deformation.

Several works [21], [31] succeeded in using the random image augmentation to improve generalization performance for semi-supervised tasks on classifying images. These methods can also be interpreted as types of techniques that smooth the model around input data points and extrapolate on the unlabeled examples. In the area of nonparametric studies, these types of methods are referred to as label propagation [44].

Another family of methods for the semi-supervised learning of NNs that are worth mentioning are those based on sophisticated generative models. These methods are different from those that we introduced above because they do not require an explicit definition of smoothness. Kingma et al. [19] applied a generative model based on a variational autoencoder to semi-supervised learning. This work was followed by several variants [24], [30]. Generative adversarial networks (GANs) proposed by [10] are a recently popular high-performance framework that can also be applied to semi-supervised learning [32], [33]. One shortcoming in the methods based on the generative model is that they are difficult to interpret. The task of analyzing "why they work" requires a series of meticulous computations and approximations that are not always justifiable. Moreover, in practice, these methods require the careful tuning of many hyperparameters in the generative model. These facts can be discouraging when considering the use of these methods in many practical applications.

3 METHOD

We begin this section with a set of notation. Let $x \in R^I$ and $y \in Q$ denote an input vector and output label, where I is the input dimension. Additionally, we denote the output distribution by parameter θ as $p(y|x,\theta)$. We use $\hat{\theta}$ to denote the vector of the model parameter during training. We use $\mathcal{D}_l = \{x_l^{(n)}, y_l^{(n)} | n = 1, \ldots, N_l\}$ to denote a labeled dataset, and $\mathcal{D}_{ul} = \{x_{ul}^{(n')} | n' = 1, \ldots, N_{ul}\}$ to denote an unlabeled dataset. We train the model $p(y|x,\theta)$ using $\mathcal{D}_l$ and $\mathcal{D}_{ul}$.

3.1 Adversarial Training

Our method is closely related to the adversarial training proposed by Goodfellow et al. [11]. We therefore formulate adversarial training before introducing our method. The loss function of adversarial training in [11] is defined as

$$L_{\mathrm{adv}}(x_l, y_l, \theta) := D\left[h(y_l), p(y|x_l + r_{\mathrm{adv}}, \theta)\right] \quad (1)$$

$$\text{where } r_{\mathrm{adv}} := \underset{r;\|r\|\leq\epsilon}{\arg\max}\, D\left[h(y_l), p(y|x_l + r, \theta)\right], \quad (2)$$

where $D[p, p']$ is a non-negative function that represents the distance between two distributions p and p'. For example, D can be the cross entropy $D[p, p'] = -\sum_i p_i \log p'_i$, where p and p' are the vectors whose i-th elements represent the probability of the i-th class. The function $h(y)$ is the distribution derived from the training sample, the function to which we strive to approximate by the parameterized model. For classification tasks, $h(y)$ is a one-hot vector of y whose entries are all zero except for the index corresponding to label y. For regression tasks, we can use the normal distribution centered at y with constant variance, or the delta function centered at y.

Generally, we cannot obtain a closed form for the exact adversarial perturbation r_{adv}. However, we can approximate r_{adv} with a linear approximation of D with respect to r in Eq.(2). When the norm is L_2, adversarial perturbation can be approximated by

$$r_{\mathrm{adv}} \approx \epsilon \frac{g}{\|g\|_2}, \text{ where } g = \nabla_{x_l} D\left[h(y), p(y|x_l, \theta)\right]. \quad (3)$$

When the norm is L_∞, adversarial perturbation can be approximated by

$$r_{\mathrm{adv}} \approx \epsilon \mathrm{sign}(g), \quad (4)$$

where g is the same function that appeared in Eq.(3). The authors of [11] originally used (4) for their adversarial training. Note that, for NNs, the gradient $\nabla_{x_l} D\left[h(y_l), p(y|x_l, \theta)\right]$ can be efficiently computed by backpropagation. Optimizing the loss function of the adversarial training in Eq.(1) based on the adversarial perturbation defined by Eq.(3)(or (4)), the authors were able to train a model with better generalization performance than the model trained with random perturbations [11], [27].

3.2 Virtual Adversarial Training

Adversarial training is a successful method that works for many supervised problems. However, rather than the one-hot vector used in the original study, a more natural choice for $h(y)$ is the true conditional probability $h(y) = q(y|x_l)$ given the input x_l. Adopting this idea, the loss function can be rewritten as

$$D\left[q(y|x_*), p(y|x_* + r_{\mathrm{qadv}}, \theta)\right]$$
$$\text{where } r_{\mathrm{qadv}} := \underset{r;\|r\|\leq\epsilon}{\arg\max}\, D\left[q(y|x_*), p(y|x_* + r, \theta)\right],$$

where x_* represents either x_l or x_{ul}. To evaluate this loss, however, we need the identity of $q(y|x)$, which remains unknown throughout the training process. Therefore, in this study, we use the current estimate $p(y|x,\hat{\theta})$ in place of $q(y|x)$. With this compromise, we arrive at our representation of Eq.(2) given by

$$\mathrm{LDS}(x_*, \theta) := D\left[p(y|x_*, \hat{\theta}), p(y|x_* + r_{\mathrm{vadv}}, \theta)\right] \quad (5)$$

$$r_{\mathrm{vadv}} := \underset{r;\|r\|_2\leq\epsilon}{\arg\max}\, D\left[p(y|x_*, \hat{\theta}), p(y|x_* + r, \theta)\right], \quad (6)$$

which defines our *virtual adversarial perturbation*. The loss $\mathrm{LDS}(x_*, \theta)$ can be considered as a measure of the local smoothness of the current model on each input data point x such that its reduction would make the model smooth at each data point. The regularization term we propose in this study is the average of $\mathrm{LDS}(x_*, \theta)$ over all input data points:

$$\mathcal{R}_{\mathrm{vadv}}(\mathcal{D}_l, \mathcal{D}_{ul}, \theta) := \frac{1}{N_l + N_{ul}} \sum_{x_* \in \mathcal{D}_l, \mathcal{D}_{ul}} \mathrm{LDS}(x_*, \theta). \quad (7)$$

The full objective function is thus given by

$$\ell(\mathcal{D}_l, \theta) + \alpha \mathcal{R}_{\mathrm{vadv}}(\mathcal{D}_l, \mathcal{D}_{ul}, \theta), \quad (8)$$

where $\ell(\mathcal{D}_l, \theta)$ is the negative log-likelihood for the labeled dataset. We refer to the training method with the regularizer $\mathcal{R}_{\mathrm{vadv}}$ as VAT.

One notable advantage of VAT is that there are just two scalar-valued hyperparameters: (1) the norm constraint for the adversarial direction $\epsilon > 0$ and (2) the regularization coefficient $\alpha > 0$ that controls the relative balance between the negative log-likelihood and the regularizer $\mathcal{R}_{\mathrm{vadv}}$. In fact, for all our experiments, our VAT achieved superior performance by tuning only the hyperparameter ϵ, while fixing $\alpha = 1$. Theoretically, these two hyperparameters play similar roles, as discussed later in Section 4.2. Thus, unlike many generative model-based methods, VAT does not suffer from the heavy computation for the optimization of numerous hyperparameters. Note that, as opposed to adversarial training [11], the definition of virtual adversarial perturbation only requires input x and does not require label y. This is the property that allows us to apply VAT to semi-supervised learning. Fig. 1 shows how VAT works on semi-supervised learning on a two-dimensional synthetic dataset. We used an NN classifier with one hidden layer with 50 hidden units. At the beginning of training, the classifier predicted different labels for input data points in the same cluster, and LDS on the boundaries was very high (see panel (2) in Fig. 1). There is a high pressure for the model to be smooth around the points with large LDS values. As training progressed, the model evolved so that the label prediction on the points with large LDS values were strongly influenced by the labeled inputs in the vicinity. This encouraged the model to predict the same label for the points that belong to the same cluster, which is what we often desire in semi-supervised learning. As we show in Fig. 1, VAT on semi-supervised learning can be given a similar interpretation as that given to label propagation [44], which is another branch of methods for semi-supervised learning.

3.3 Fast Approximation Method for r_{vadv} and the Derivative of the Objective Function

Once virtual adversarial perturbation r_{vadv} is computed, the evaluation of $\mathrm{LDS}(x_*, \theta)$ simply becomes the computation of the divergence D between the output distributions $p(y|x_*, \hat{\theta})$ and $p(y|x_* + r_{\mathrm{vadv}}, \theta)$. However, the evaluation of r_{vadv} cannot be performed with the linear approximation as in the original adversarial training in Eq.(3) because the gradient of $D[p(y|x_*, \hat{\theta}), p(y|x_* + r, \hat{\theta})]$ with respect to r on $r = 0$ is always 0 for current parameter $\hat{\theta}$. In the following, we propose an efficient computation of r_{vadv}, for which there is no evident closed form.

For simplicity, we denote $D[p(y|x_*, \hat{\theta}), p(y|x_* + r, \theta)]$ by $D(r, x_*, \theta)$. We assume that $p(y|x_*, \theta)$ is twice differentiable with respect to θ and x almost everywhere. Because $D(r, x_*, \hat{\theta})$ takes the minimal value at $r = 0$, the differentiability assumption dictates that its first derivative $\nabla_r D(r, x, \hat{\theta})|_{r=0}$ is zero. Therefore, the second-order Taylor approximation is given by

$$D(r, x, \hat{\theta}) \approx \frac{1}{2} r^T H(x, \hat{\theta}) r, \tag{9}$$

where $H(x, \hat{\theta})$ is the Hessian matrix given by $H(x, \hat{\theta}) := \nabla\nabla_r D(r, x, \hat{\theta})|_{r=0}$. Under this approximation, r_{vadv} emerges as the first dominant eigenvector $u(x, \hat{\theta})$ of $H(x, \hat{\theta})$ with magnitude ϵ:

$$\begin{aligned} r_{\mathrm{vadv}} &\approx \arg\max_r \{r^T H(x, \hat{\theta}) r;\ \|r\|_2 \le \epsilon\} \\ &= \overline{\epsilon u(x, \hat{\theta})}, \end{aligned} \tag{10}$$

where $\bar{v}$ denotes a unit vector whose direction is the same as its argument vector v; that is, $\bar{v} \equiv \frac{v}{\|v\|_2}$. Hereafter, we denote $H(x, \hat{\theta})$ by H for simplicity.

The next challenge is to address the $O(I^3)$ runtime required for the computation of the eigenvectors of the Hessian H. This can be a large burden for high-dimensional problems. We resolve this issue with the approximation via a power iteration method [8] and finite difference method. Let d be a randomly sampled unit vector. Provided d is not perpendicular to the dominant eigenvector u, the iterative calculation of

$$d \leftarrow \overline{Hd} \tag{11}$$

makes d converge to u. To reduce the computational time, we perform this operation without the direct computation of H. Note that Hd can be approximated using a finite difference method:

$$\begin{aligned} Hd &\approx \frac{\nabla_r D(r, x, \hat{\theta})|_{r=\xi d} - \nabla_r D(r, x, \hat{\theta})|_{r=0}}{\xi} \\ &= \frac{\nabla_r D(r, x, \hat{\theta})|_{r=\xi d}}{\xi}, \end{aligned} \tag{12}$$

with $\xi \neq 0$. In the computation above, we use the fact that $\nabla_r D(r, x, \hat{\theta})|_{r=0} = 0$ again. To summarize, we can approximate r_{vadv} with the repeated application of the following update:

$$d \leftarrow \overline{\nabla_r D(r, x, \hat{\theta})|_{r=\xi d}}. \tag{13}$$

The computation of $\nabla_r D$ can be performed in a straightforward manner. For NNs, this can be performed with one set of backpropagation. The approximation introduced here improves monotonically with increasing K, which is the number of power iterations. Thus, for NNs, the computation of r_{vadv} can be performed with K sets of backpropagation. Surprisingly, only one power iteration was sufficient for high performance on various benchmark datasets. This approximation of r_{vadv} with $K = 1$ results in an approximation that is similar in form to Eq.(3):

$$r_{\mathrm{vadv}} \approx \epsilon \frac{g}{\|g\|_2} \tag{14}$$

$$\text{where } g = \nabla_r D\left[p(y|x, \hat{\theta}), p(y|x + r, \hat{\theta})\right]\Big|_{r=\xi d}. \tag{15}$$

We further discuss the effects of the number of power iterations in Section 4.

After computing r_{vadv}, the derivative of $\mathcal{R}_{\mathrm{vadv}}$ can be easily computed with one set of forward and backpropagation on NNs. By contrast, the derivative of r_{vadv} is not only convoluted and computationally costly, but also introduces another source of variance to the gradient, which exacerbates the performance of the algorithm. Our VAT therefore ignores the dependency of r_{vadv} on θ. In total, including the log-likelihood term, the derivative of the full objective function (8) can be computed with $K + 2$ sets of

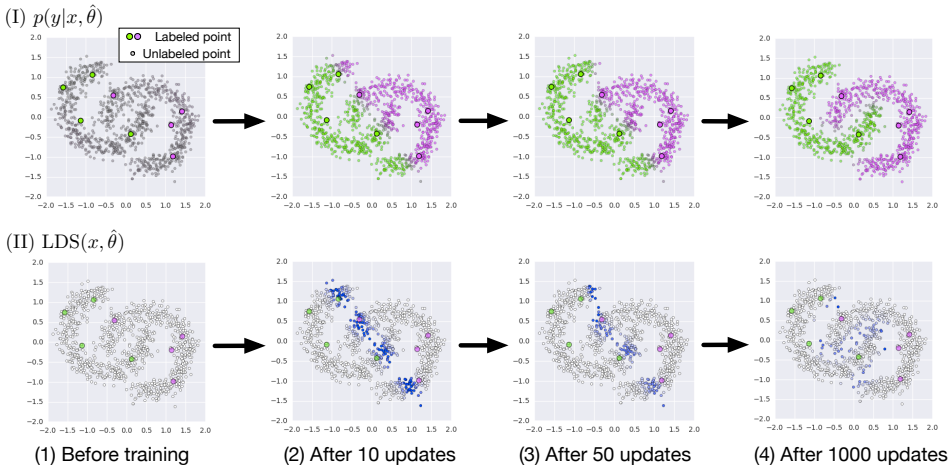

Fig. 1: Demonstration of how our VAT works on semi-supervised learning. We generated 8 labeled data points ($y = 1$ and $y = 0$ are green and purple, respectively), and 1,000 unlabeled data points in 2-D space. The panels in the first row (I) show the prediction $p(y = 1|x, \theta)$ on the unlabeled input points at different stages of the algorithm (green, gray, and purple, which correspond to the values 1.0, 0.5, and 0.0, respectively). The panels in the second row (II) are heat maps of the regularization term LDS$(x, \hat{\theta})$ on the input points. The values of LDS on blue-colored points are relatively high in comparison to the gray-colored points. We used KL divergence as D in Eq.(5). Note that, at the onset of training, all the data points contributed to the classifier in a similar manner. After 10 updates, the model boundary was still appearing *over* the inputs. As training progressed, VAT pushed the boundary away from the labeled input data points.

backpropagation. Algorithm 1 summarizes the procedure for calculating $\nabla_\theta \mathcal{R}_{\text{vadv}}$ with mini-batch SGD with one power iteration. VAT is an algorithm that updates the model by the weighted sum of the gradient of the likelihood and $\nabla_\theta \mathcal{R}_{\text{vadv}}$ computed by Algorithm 1.

Algorithm 1 Mini-batch SGD for $\nabla_\theta \mathcal{R}_{\text{vadv}}(\theta)|_{\theta=\hat{\theta}}$, with a one-time power iteration method.

1) Choose M samples of $x^{(i)} (i = 1, \ldots, M)$ from dataset $\mathcal{D}$ at random.
2) Generate a random unit vector $d^{(i)} \in R^I$ using an iid Gaussian distribution.
3) Calculate r_{vadv} via taking the gradient of D with respect to r on $r = \xi d^{(i)}$ on each input data point $x^{(i)}$:
$$g^{(i)} \leftarrow \nabla_r D \left[p(y|x^{(i)}, \hat{\theta}), p(y|x^{(i)} + r, \hat{\theta}) \right] \Big|_{r=\xi d^{(i)}},$$
$$r^{(i)}_{\text{vadv}} \leftarrow g^{(i)} / \|g^{(i)}\|_2$$
4) **Return**
$$\nabla_\theta \left(\frac{1}{M} \sum_{i=1}^{M} D \left[p(y|x^{(i)}, \hat{\theta}), p(y|x^{(i)} + \tilde{r}^{(i)}_{\text{vadv}}, \theta) \right] \right) \Big|_{\theta=\hat{\theta}}$$

3.4 Virtual Adversarial Training vs. Random Perturbation Training

We define a regularization function given by

$$\mathcal{R}^{(K)}(\theta, \mathcal{D}_l, \mathcal{D}_{ul}) := \frac{1}{N_l + N_{ul}} \sum_{x \in \mathcal{D}_l, \mathcal{D}_{ul}} E_{r_K} \left[D \left[p(y|x, \hat{\theta}), p(y|x + r_K, \theta) \right] \right], \tag{16}$$

where r_K is sampled from the uniform distribution on the sphere $U(r|\epsilon)$ with radius ϵ followed by the aforementioned K-times power iterations. In practice, we use an empirical expectation about random perturbation r_K. When $K \geq 1$, this regularizer coincides with the implementation of VAT. We refer to training with $\mathcal{R}^{(0)}$ as RPT. This is a simpler version of VAT that does not perform power iterations, and we devised the name from the fact that it only smooths the function isotropically around each input data point.

As we discuss further in Section 4, RPT falls behind VAT in its sheer ability to reduce the generalization error. There could be two reasons for the superiority of VAT. First, its regularization functions have essentially different objectives. For each observed input point, VAT trains the model to assign the same label distribution only to the set of proximal points aligned in the adversarial direction. By contrast, RPT encourages the model to assign the same label distribution to all input points in the neighborhood of each observed input. Second, the learning process of VAT is inherently more stable than that of RPT. At each step of the algorithm, the power iteration generates a vector that has a large projection to the virtual adversarial direction with high probability. Note that, as $K \to \infty$, under the sufficient regularity of the model, the gradient of the expression (16) approaches a deterministic value $\frac{\partial}{\partial \theta} \lambda_1(x, \theta)$, where λ_1 is the dominant eigenvector of $H(x, \theta)$. Thus, the direction to which VAT smooths the model is more deterministic than the direction to which RPT smooths the model, which is uniformly distributed over the sphere of radius ϵ; the stability of the learning of RPT always suffers from the variance of (16). In Section 4.5, we investigate the variance of the gradients in more detail and compare RPT and VAT from this perspective.

4 EXPERIMENTS

We conducted a set of numerical experiments to assess the following aspects of VAT:

- the sheer efficacy of VAT in comparison with RPT and recent collection of competitive algorithms for supervised and semi-supervised learning,
- the effect of hyperparameters (the perturbation size ϵ, the regularization coefficient α, and the number of power iterations) on the performance of VAT,
- VAT's effect on the robustness of the trained NNs against virtual adversarial examples, and
- the mechanism behind the advantage of using virtual adversarial perturbation as opposed to random perturbation.

In the following, we describe the experimental settings and outcomes of the experiments. For more details on the dataset and model architectures, please see Appendix . We used Theano [38] and TensorFlow [1] to train the NN and CNN models[1]. Throughout our experiments, we used fully connected NNs or convolutional neural networks (CNNs) as the architectures of classifiers. We use $p(y|x, \theta)$ to denote the label distribution of the classifier, where θ represents the parameters of the NN. For the activation function, we used ReLU [7], [17], [28]. We also used Batch Normalization [16]. For the divergence D in Eq. (1), we chose KL divergence:

$$D\left[p(y|x), p(y|x+r)\right] := \sum_{y \in Q} p(y|x) \log \frac{p(y|x)}{p(y|x+r)}, \quad (17)$$

where Q is the domain of y. For classification problems, it is the set of all possible labels. We set ξ in Eq.(12) to be $1e$-6 in all of experiments.

4.1 Testing the Efficacy of VAT on Benchmark Tasks

4.1.1 Supervised Learning on MNIST and CIFAR-10

We tested the performance of VAT on the MNIST dataset. We used an NN with four hidden layers, whose numbers of units were $(1200, 600, 300, 150)$. We provide more details of the experiments in Appendix A.

For each regularization method, we used the set of hyperparameters that achieved the best performance on the validation dataset of size $10,000$, which was selected from the pool of training samples of size $60,000$. We also prepared a test dataset that had no intersection with the training dataset. We input the test dataset into the trained NNs and recorded their test errors. We repeated this procedure 10 times with different random seeds for the weight initialization, and reported the average test errors.

Fig. 2 shows the transition of $\mathcal{R}_{\mathrm{vadv}}$ and the learning curves for the baseline NN trained *without* VAT (denoted by 'wo/ VAT') and the NN trained *with* VAT (denoted by 'w/ VAT'). At the beginning of training, $\mathcal{R}_{\mathrm{vadv}}$ on both the baseline and VAT were 0 because the network weights were initialized randomly. As training progressed, $\mathcal{R}_{\mathrm{vadv}}$ of the NN trained with VAT exceeded that of the baseline; that is, the model trained with VAT grew smoother than the baseline in terms of LDS.

1. The Tensorflow implementation for the experiments is available at https://github.com/takerum/vat_tf. Chainer [40] implementation is also available at https://github.com/takerum/vat_chainer.

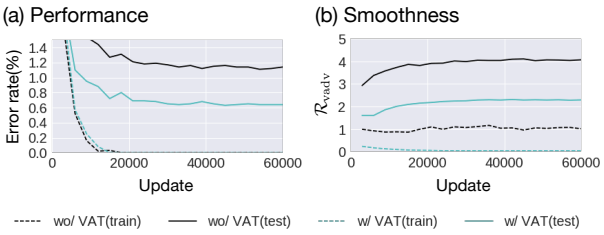

Fig. 2: Transition of the (a) classification error and (b) $\mathcal{R}_{\mathrm{vadv}}$ for supervised learning on MNIST. We set $\epsilon = 2.0$ for the evaluation of $\mathcal{R}_{\mathrm{vadv}}$ on both the baseline and VAT. This is the value with which the VAT-trained model achieved the best performance on the validation dataset.

Table 1 summarizes the test errors achieved by our regularization method (VAT) and the other regularization methods. VAT performed better than all the contemporary methods except ladder networks, which are a highly advanced method based on the generative model.

We also tested VAT with $K > 1$. As we discuss in Section 4.3, however, we were not able to confirm a substantial improvement over $K = 1$.

We also ran our algorithm on CIFAR-10 [20]. For this set of experiments, we used a 'Conv-Large' in Table 6 in Appendix D with a dropout, as the baseline. Table 2 lists the performance of the close family of CNN-based classifiers that do not implement recent advanced architectures, such as ResNet [14] and DenseNet [15]. For CIFAR-10, our VAT achieved satisfactory performance relative to the standards.

TABLE 1: Test performance of supervised learning on MNIST with 60,000 labeled examples in the permutation invariant setting. The first row cites the results provided by the original paper. The second row shows the performance achieved by the authors' implementation.

Method	Test error rate(%)
SVM (Gaussian kernel)	1.40
Dropout [35]	1.05
Adversarial, L_∞ norm constraint [11]	0.78
Ladder networks [30]	**0.57**
Baseline (MLE)	1.11
RPT	0.82
Adversarial, L_∞ norm constraint	0.79
Adversarial, L_2 norm constraint	0.71
VAT	0.64

TABLE 2: Test performance of CNNs on the supervised CIFAR-10 with 50000 labeled examples.

Method	Test error rate(%)
Network in Network [23]	8.81
All-CNN [34]	7.25
Deeply Supervised Net [22]	7.97
Highway Network [36]	7.72
Baseline (only with dropout)	6.76
RPT	6.25
VAT	**5.81**

4.1.2 Semi-Supervised Learning on MNIST, SVHN, and CIFAR-10

Recall that our definition of $\mathrm{LDS}(x, \theta)$ at any point x is independent of its label information y. In particular, this means that we can apply VAT to semi-supervised learning tasks. We emphasize that this is a property not shared by adversarial training. We applied VAT to semi-supervised learning tasks on three image classification datasets: MNIST, SVHN [29], and CIFAR-10 [20]. We provide the details of the experimental settings in Appendix D.

For MNIST, we used the same architecture of NN as in the previous section. We also used normalization in our implementation. We used a mini-batch of size 64 for the calculation of the negative log-likelihood term, and a mini-batch of size 256 for the calculation of $\mathcal{R}_{\mathrm{vadv}}$ in Eq. (8). As mentioned in Section 3.2, we used both labeled and unlabeled sets for the calculation of $\mathcal{R}_{\mathrm{vadv}}$. Table 3 summarizes the results for the permutation invariant MNIST task. All the methods belong to the family of semi-supervised learning methods. For the MNIST dataset, VAT outperformed all the contemporary methods other than the methods based on generative models, such as [24], [30], [32], which are current state-of-the-art methods.

TABLE 3: Test performance on the semi-supervised MNIST with 100 and 1,000 labeled examples on the permutation invariant setting.

Method	Test error rate(%) $N_l = 100$	$N_l = 1000$
TSVM [5]	16.81	5.38
Pseudo Ensembles Agreement (PEA) [3]	5.21	2.87
Deep Generative Model [19]	3.33	2.59
CatGAN [33]	1.91	1.73
Ladder Networks [30]	1.06	**0.84**
GAN with feature matching [32]	**0.93**	
RPT	6.81	1.58
VAT	1.36	1.27

For the experiments on SVHN and CIFAR-10, we used two types of CNNs (Conv-Small and Conv-Large) used in recent state-of-the-art semi-supervised learning methods [21], [30], [32], [33]. 'Conv-Small CNN' has practically the same structure as the CNN used in [32], and 'Conv-Large CNN' has practically the same structure as the CNN used in [21]. We provide more details of the architectures in Table 6 in Appendix D.

Also, In SVHN and CIFAR-10 experiments, we adopted conditional entropy of $p(y|x, \theta)$ as a additional cost, used in [12] and similar one used in [31]:

$$\begin{aligned}\mathcal{R}_{\mathrm{cent}} &= H(Y|X) \\ &= -\frac{1}{N_l + N_{ul}} \sum_{x \in \mathcal{D}_l, \mathcal{D}_{ul}} \sum_y p(y|x, \theta) \log p(y|x, \theta).\end{aligned} \tag{18}$$

The conditional entropy minimization forces the model $p(y|x, \theta)$ to exaggerate the prediction on each data point. For semi-supervised image classification tasks, this additional cost is especially helpful because it ensures that one category is assigned to each image with high probability. In what follows, 'VAT+EntMin' indicates the training with $\mathcal{R}_{\mathrm{vadv}} + \mathcal{R}_{\mathrm{cent}}$.

Table 4 summarizes the results for SVHN and CIFAR-10. Our method achieved the test error rate of 14.82(%) with VAT, which outperformed the current state-of-the-art methods for semi-supervised learning on CIFAR-10. VAT+EntMin outperformed the current state-of-the-art methods for semi-supervised learning on both SVHN and CIFAR-10.

TABLE 4: Test performance on SVHN (1,000 labeled) and CIFAR-10 (4,000 labeled) without image data augmentation.

Method	Test error rate(%) SVHN $N_l = 1000$	CIFAR-10 $N_l = 4000$
SWWAE [43]	23.56	
Skip Generative Model [24]	16.30	
Ladder networks, Γ model [30]		20.40
CatGAN [33]		19.58
GAN with feature matching [32]	8.11	18.63
$\prod$ model [21]	5.43	16.55
(on Conv-Small used in [32])		
RPT	8.41	18.56
VAT	**6.83**	**14.87**
(on Conv-Large used in [21])		
VAT	5.77	14.82
VAT+EntMin	**4.28**	**13.15**

Table 5 shows the performance of VAT and contemporary methods implemented with moderate image data augmentation (translation and horizontal flip). All methods other than VAT were implemented with the additional knowledge that the label of the image did not change with augmentation. On CIFAR-10, VAT still outperformed the listed methods with this handicap. VAT+EntMin with moderate data augmentation also outperformed the current state-of-the-art methods for semi-supervised learning on both SVHN and CIFAR-10.

TABLE 5: Test performance on SVHN and CIFAR-10 with image data augmentation. The performance of all methods other than Sajjadi et al. [31] are based on experiments with the moderate data augmentation of translation and flipping (see Appendix D for more detail). Sajjadi et al. [31] used extensive image augmentation, which included rotations, stretching, and shearing operations.

Method	Test error rate(%) SVHN $N_l = 1000$	CIFAR-10 $N_l = 4000$
Γ model [30] (experimented by [21])		≈ 16
$\prod$ model [21].	4.84	12.36
Temporal ensembling [21]	4.43	12.16
Sajjadi et al. [31]		11.29
(On Conv-Large used in [21])		
VAT	5.42	11.36
VAT+EntMin	**3.86**	**10.55**

4.2 Effects of Perturbation Size ϵ and Regularization Coefficient α

Generally, there is a high computational cost for optimizing the hyperparameters of large NNs on a large dataset. One advantage of VAT is that the algorithm involves only two hyperparameters: α and ϵ. Even better, our experiments

suggest that, in an appropriate setting, VAT works with the optimization of ϵ alone. In all our experiments in Section 4.1, VAT achieved competitive results with a fixed value of $\alpha = 1$.

This treatment is not entirely baseless. Note that, for small ϵ, the hyperparameter α plays a similar role to ϵ. To see this, consider the Taylor expansion of Eq.(5) for small ϵ, given by

$$\max_r\{D(r,x,\theta); \|r\|_2 \le \epsilon\} \approx \max_r\{\frac{1}{2} r^T H(x,\theta) r; \|r\|_2 \le \epsilon\} = \frac{1}{2}\epsilon^2 \lambda_1(x,\theta), \quad (19)$$

where $H(x,\theta)$ and $\lambda_1(x,\theta)$ are the Hessian matrix of D in Eq.(9) and its dominant eigenvalue, respectively. Substituting this into the total objective function (Eq.(8)), we obtain

$$\begin{aligned} &\ell(\theta, \mathcal{D}_l) + \alpha \mathcal{R}_{\mathrm{vadv}}(\theta, \mathcal{D}_l, \mathcal{D}_{ul},) \\ =&\ell(\theta, \mathcal{D}_l) + \alpha \frac{1}{N_l + N_{ul}} \sum_{x_* \in \mathcal{D}_l, \mathcal{D}_{ul}} \max_r \{D(r, x_*, \theta), \|r\|_2 \le \epsilon\} \\ \approx&\ell(\theta, \mathcal{D}_l) + \frac{1}{2}\alpha\epsilon^2 \frac{1}{N_l + N_{ul}} \sum_{x_* \in \mathcal{D}_l, \mathcal{D}_{ul}} \lambda_1(x_*, \theta). \end{aligned} \quad (20)$$

Thus, at least for small ϵ, the strength of the regularization in VAT is proportional to the product of the two hyperparameters, α and ϵ^2; that is, in the region of small ϵ, the hyperparameter search for either ϵ or α suffices. However, when we consider a relatively large value of ϵ, the hyperparameters α and ϵ cannot be brought together. In such a case, we strive to search for the best pair of hyperparameters that attains optimal performance.

In our experiments on MNIST, tuning ϵ alone sufficed for achieving satisfactory performance. On an empirical basis, we therefore recommend that the user prioritizes the parameter search for ϵ over the search for α.

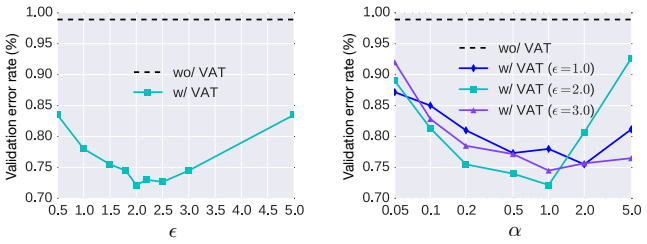

(a) Effect of ϵ ($\alpha = 1$). (b) Effect of α.

Fig. 3: Effect of ϵ and α on the validation performance for supervised MNIST.

Fig. 3 shows the effects of ϵ and α on the validation performance of supervised learning on MNIST. In Fig. 3a, we show the effect of ϵ with fixed $\alpha = 1$, and in Fig. 3b we show the effects of α with different values of ϵ in the range $\{1.0, 2.0, 3.0\}$. During the parameter search of ϵ for supervised learning on MNIST, the algorithm performed optimally when $\alpha = 1$. Based on this result, we fixed $\alpha = 1$ while searching for optimal ϵ in all other experiments, including both supervised learning on MNIST and CIFAR-10, as well as unsupervised learning on MNIST, CIFAR-10, and SVHN. As we showed in Section 4.1, this simple tuning of ϵ was sufficient for the good performance of VAT, and it even achieved state-of-the-art performance for several tasks.

4.3 Effect of the Number of Power Iterations K

Fig. 4 shows the $\mathcal{R}_{\mathrm{vadv}}$ value of the models trained with different K (the number of power iterations) for supervised learning on MNIST and semi-supervised learning on CIFAR-10. We can observe a significant increase in the value of $\mathcal{R}_{\mathrm{vadv}}$ over the transition from $K = 0$ (random perturbations) to $K = 1$ (virtual adversarial perturbations). We can also observe that the value saturates at $K = 1$. In fact, we did not observe much improvement in performance with larger values of K.

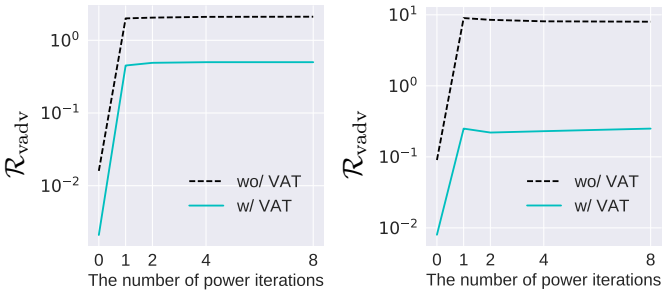

(a) MNIST (b) CIFAR-10

Fig. 4: Effect of the number of power iterations on $\mathcal{R}_{\mathrm{vadv}}$ for (a) supervised MNIST and (b) semi-supervised CIFAR-10.

4.4 Visualization of Virtual Adversarial Examples

4.4.1 *Virtual Adversarial Examples Produced by the Model Trained with Different Choices of ϵ*

One might be interested in the actual visual appearance of the virtual adversarial examples with an appropriate size of ϵ that improves the regularization performance. In Fig. 5, we aligned (I) the transition of the performance of VAT on SVHN and CIFAR-10 with respect to ϵ along (II) the actual virtual adversarial examples that the corresponding models generated at the end of training. For small ϵ (designated as (1) in the figure), it is difficult for human eyes to distinguish the virtual adversarial examples from the clean images. The size of ϵ that achieved the best validation performance is designated as (2). The virtual adversarial perturbations with ϵ of size (2) are on the verge of the level of corruption that can alter human perception. For a larger value of ϵ (designated as (3) in the figure), we can clearly observe the effect of the over-regularization of the model. In fact, the virtual adversarial examples generated by the models trained with this range of ϵ are very far from the clean image, and we observe that the algorithm did the uncessary work of smoothing the label distribution over the set of images that are "unnatural."

4.4.2 *Robustness against Virtual Adversarial Examples after Training*

We studied the nature of the robustness that our VAT achieved during training. We trained CNNs on CIFAR-10 with and without VAT and prepared pairs of virtual

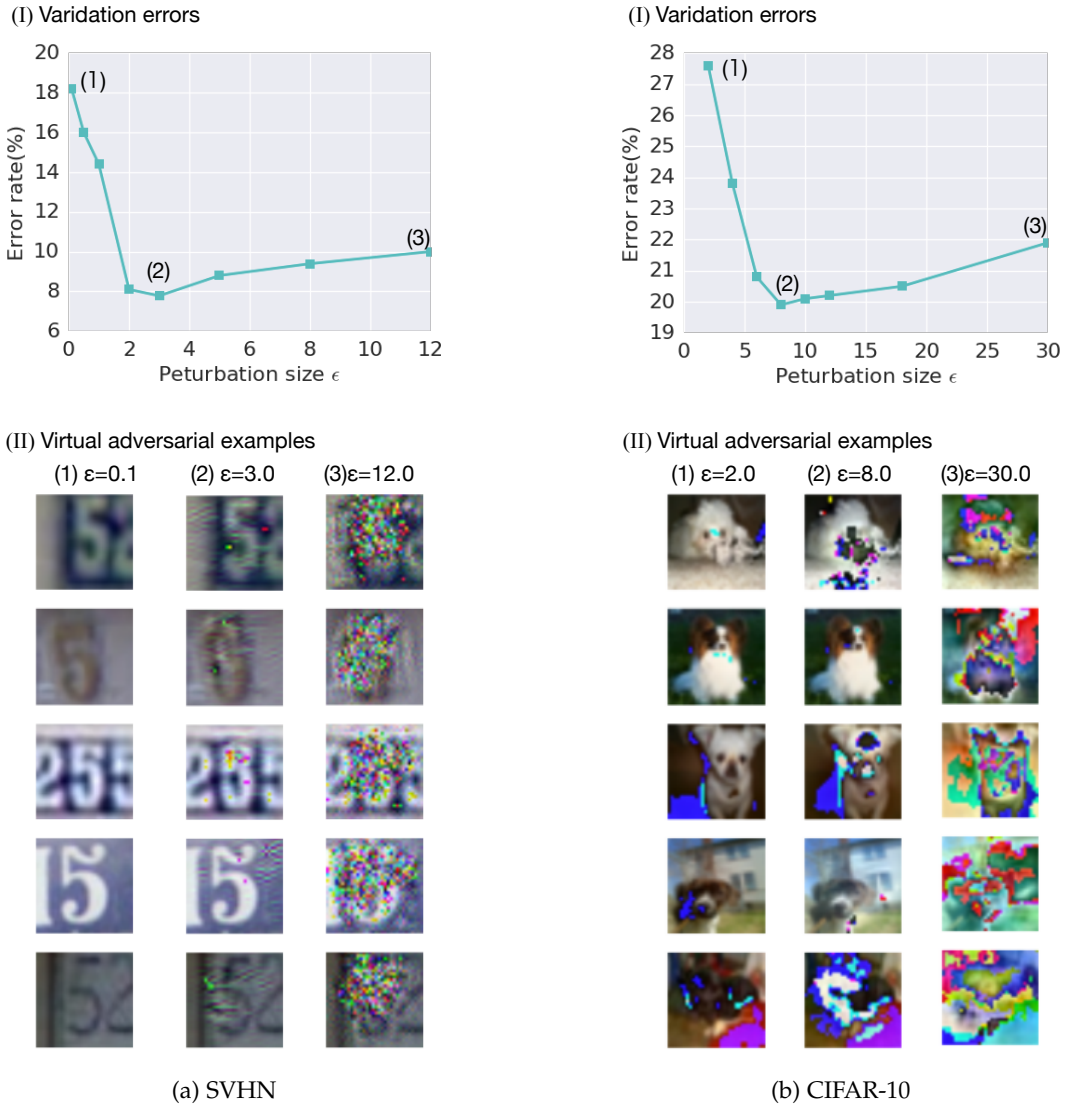

(a) SVHN (b) CIFAR-10

Fig. 5: **Performance of VAT with different values of ϵ.** The effect of ϵ on the performance of semi-supervised learning, together with the set of typical virtual adversarial examples generated by a VAT-trained model with the corresponding value of ϵ.

adversarial examples generated from the same set of pictures, each consisting of (1) the virtual adversarial example generated by the model trained with VAT and (2) the virtual adversarial example generated by the model trained without VAT. We studied the rates of misidentification on these pairs of adversarial examples made by the models trained with and without VAT.

Fig. 6 shows the rates at which the two models misidentified the images perturbed with virtual adversarial perturbations of different sizes. The first column (A) in the middle panel corresponds to the rates of misidentification made on the virtual adversarial examples generated by the model trained with VAT. The second column (B) corresponds to the rates on the virtual adversarial examples generated by the model trained without VAT. The example pictures shown beneath the graphs are the adversarial examples generated from the set of images on which both the model trained with VAT and the model trained without VAT made a correct identification. As expected, the error rates increased monotonically with the intensity of corruption for both the model trained with VAT and model trained without VAT. Overall, we recognize that the adversarial examples generated by both models are identical to the original image in human eyes when $\epsilon \sim 10^{-1}$. The adversarial examples around $\epsilon \sim 10^{0}$ seem far from the original , but somewhat recognizable by human eyes. The virtual adversarial examples with this range of ϵ are therefore the examples on which we wish the classifier to make no mistakes. We can clearly observe that the rate of misidentification made by the VAT-trained model is much lower than that of the model trained without VAT. In particular, note in the bottom panel that the model trained with VAT correctly identifies both the adversarial examples generated by itself and the adversarial examples generated by the model trained without VAT for this range of ϵ. Simultaneously, note that the model trained with VAT alters its decision on the images when the perturbation is

too large to be recognizable. By contrast, the model trained without VAT is assigning the same labels to these over-perturbed images with a much higher rate than the VAT-trained model, which is completely unnecessary in practice. Thus, we observe that the VAT-trained model behaves much more 'naturally' than the model trained without VAT.

4.5 Experimental Assessment of the Difference between VAT and RPT

As we showed in the previous section, VAT outperforms RPT. There are two possible reasons for this difference in performance. First, as mentioned in Section 3.4, the power iteration in VAT promotes a *faster learning process* by decreasing the variance of the derivative of the objective function. Second, smoothing the function in the direction to which the model is *most sensitive* seems to be much more effective in improving generalization performance than RPT, which smooths the model isotropically around the input. We discuss the extent to which these claims are true. Let $\mathcal{D}_M$ be the mini-batch of size M randomly extracted from $\mathcal{D}$ in the algorithm, and let $\hat{\mathcal{R}}^{(K)}(\theta; \mathcal{D}_M, r_K)$ be the approximation of $\mathcal{R}^{(K)}(\theta, \mathcal{D})$ computed with the mini-batch $\mathcal{D}_M$ and a random perturbation r_K generated by K-times power iterations. To quantify the magnitude of the variance of the stochastic gradient, we define the *normalized standard deviation (SD) norm* for the gradient $\nabla_\theta \hat{\mathcal{R}}^{(K)}(\theta; \mathcal{D}_M, r_K)$ of the regularization term $\hat{\mathcal{R}}^{(K)}(\theta; \mathcal{D}_M, r_K)$, which is given as the square root of the trace of its variance normalized by the L_2 norm of its expectation:

$$\begin{aligned}&\text{(normalized SD norm)}\\ &\equiv \frac{\sqrt{\text{trace}\left(\text{Var}_{\mathcal{D}_M, r_K}\left[\nabla_\theta \hat{\mathcal{R}}^{(K)}(\theta; \mathcal{D}_M, r_K)\right]\right)}}{\left\| E_{\mathcal{D}_M, r_K}\left[\nabla_\theta \hat{\mathcal{R}}^{(K)}(\theta; \mathcal{D}_M, r_K)\right]\right\|_2}, \end{aligned} \tag{21}$$

where $\text{Var}_{\mathcal{D}_M, r_K}[\cdot]$ and $E_{\mathcal{D}_M, r_K}[\cdot]$ represent the variance and expectation, respectively, with respect to the randomly selected mini-batch $\mathcal{D}_M$ and perturbation r_K, respectively. Fig. 7 shows the transition of the normalized SD norm during VAT of NNs for the supervised learning task on MNIST (Fig. 7a) and the semi-supervised learning task on CIFAR-10 (Fig. 7b) computed with $K = 0$ and $K = 1$ (i.e., RPT and VAT). We set M to be 100 and 128 on MNIST and CIFAR-10, which is same as the values used in training. From the figure, we can observe that the normalized SD norm for $K = 1$ is lower than that for $K = 0$ in most of the early stages of training for both MNIST and CIFAR-10. Even for the instances over which the normalized SD norm of the gradient of RPT falls below that of VAT, the difference is not substantially large. Additionally, for MNIST, the training phase over which the normalized SD norm of the gradient of VAT stays below that of the RPT coincides with the phase over which the network's $\mathcal{R}_{\text{vadv}}$ grows.

For MNIST, the normalized SD norm of RPT becomes as large as 3 times that of VAT. To understand how much the normalized SD norm affects performance, we compared RPT with the 9 sampled perturbations and optimized α against VAT with 1 sampled perturbation and $\alpha = 1$. With this setting, the normalized SD norm of RPT is not greater than that of VAT at the beginning of training (Fig. 8(a)). Note that the normalized SD norm of the gradient does not depend on α. Remarkably, even with optimal α and increased S, the performance of RPT still falls behind that of VAT with $\alpha = 1$. Additionally, we found that, even with same size of perturbation considered in the procedures, the model trained with VAT is more robust than the model trained with RPT against virtual adversarial perturbation (Fig. 8(b)). This, in particular, means that we cannot explain the superiority of VAT over RPT by the reduction in the variance alone. All these results suggest that the superior performance of VAT owes much to the unique nature of its objective function. Intuition tells us that the "max" component in the objective function as opposed to "expectation" is working to the advantage of performance. The model with arbitrary parameters lack isotropic smoothness around the sample input points; it is natural that, to fix this, we must endeavor to smooth the model *anisotropically* in the direction in which the model is *most* sensitive; that is, we need to take the maximum.

5 Conclusion

The results of our experiments on the three benchmark datasets, MNIST, SVHN, and CIFAR-10 indicate that VAT is an effective method for both supervised and semi-supervised learning. For the MNIST dataset, VAT outperformed recent popular methods other than ladder networks, which are the current state-of-the-art method based on the generative model. VAT also greatly outperformed the current state-of-the-art semi-supervised learning methods for SVHN and CIFAR-10.

The simplicity of our method is also worth re-emphasizing. With our approximation of $\mathcal{R}_{\text{vadv}}$, VAT can avoid an internal optimization when choosing the adversarial direction and can be implemented with relatively small computational cost. Additionally, VAT has only two hyperparameters (ϵ and α) and worked sufficiently well on the benchmark dataset with the optimization of only one hyperparameter: ϵ. This is a clear advantage over the methods based on generative models, whose performance is often highly dependent on computationally heavy, careful tuning of many hyperparameters.

Simultaneously, we confirmed the superiority of generative model-based methods on several experimental settings [30], [32]. Essentially, most generative model-based methods promote generalization performance by making the model robust against perturbations in the region of high $p(x)$ values, the region over which we are likely to receive new input data points in the future. In principle, this is complementary to our method, which aims to isotropically smooth the output distribution $p(y|x)$ over each observed input point without any assumption on the input distribution $p(x)$. In future work, we will consider developing a theory that can combine these two perspectives.

Acknowledgments

This study was supported by the New Energy and Industrial Technology Development Organization (NEDO), Japan.

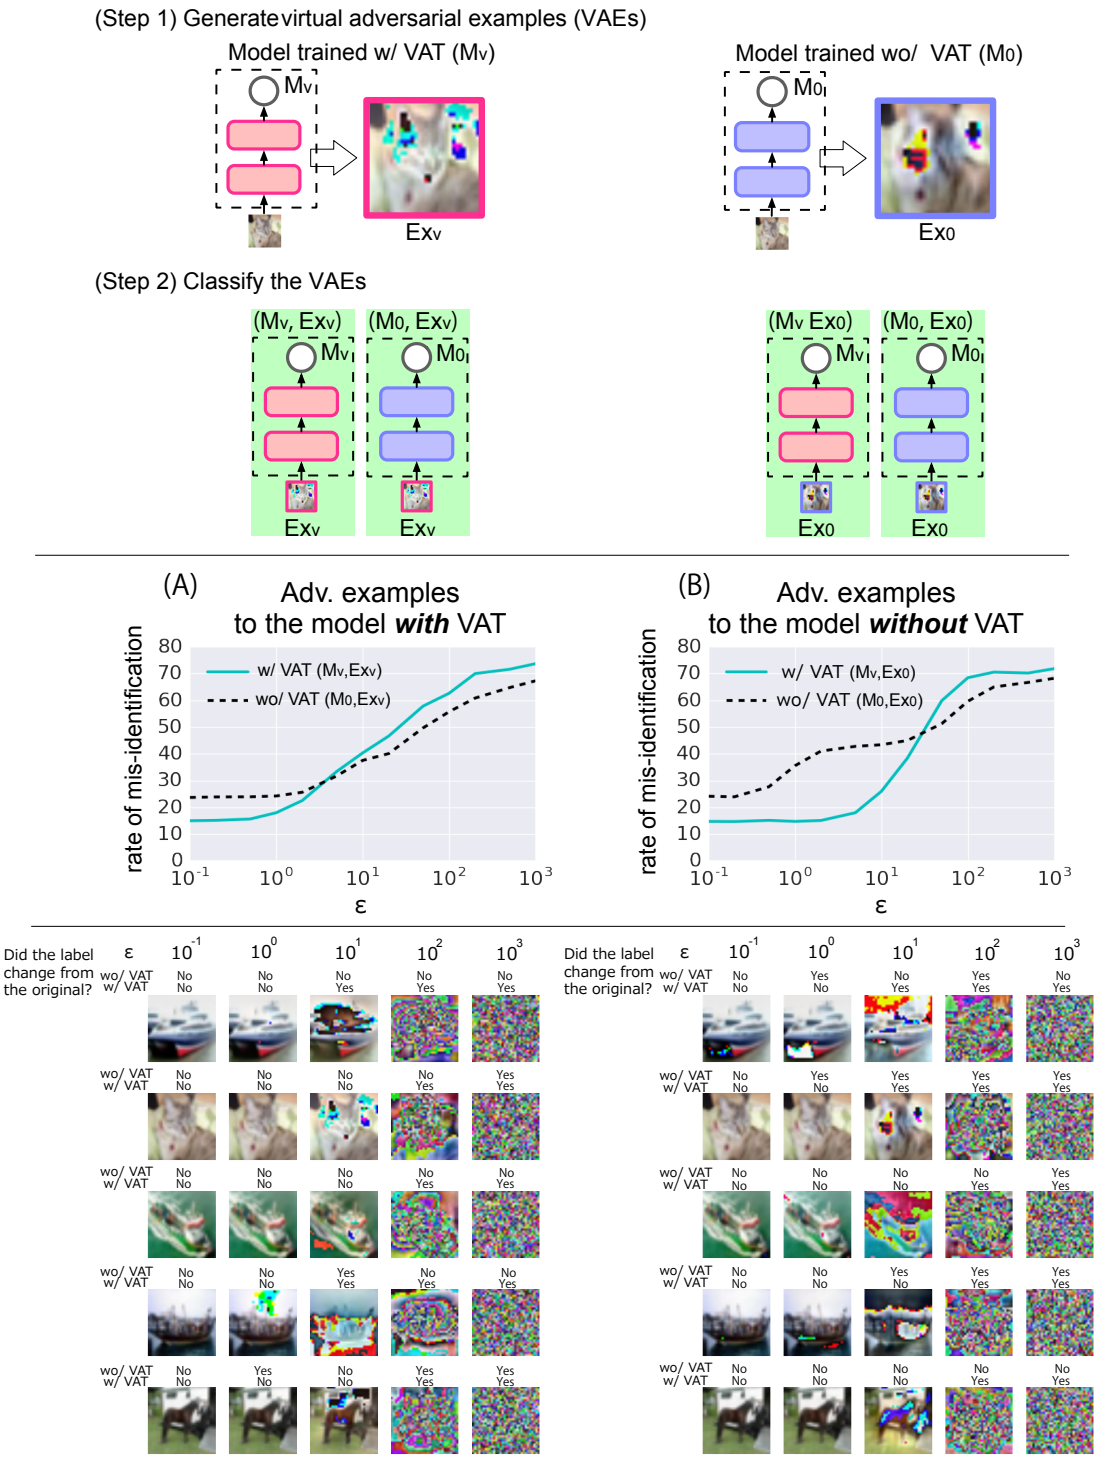

Fig. 6: **Robustness of the VAT-trained model against perturbed images.** The upper panel shows the procedure for the evaluation of robustness. In step 1, we prepared two classifiers – one trained with VAT (M_v) and another trained without VAT (M_0) – and generated a virtual adversarial example from each classifier (Ex_v and Ex_0). In step 2, we classified Ex_v and Ex_0 to these two models, thereby yielding a total of four classification results. The middle panel (graph A and B) plots the rate of misidentification made in these four classification tasks against the size of the perturbation (ϵ) used to generate the virtual adversarial examples (VAEs) in step 1. The left half of the bottom panel aligned with the graph (A) shows a set of Ex_v generated with different values of ϵ, together with the results of M_v and M_0 on the images. All Ex_v listed are images generated from the clean examples on which both M_v and M_0 made a *correct* identification. The right half of the bottom panel aligned with graph (B) shows the set of Ex_0 generated with different values of ϵ. The label 'YES' indicates that the model changed the label assignment when the perturbation was applied to the image. The label 'No' indicates that the model maintained the label assignment on the perturbed image. Note that the M_v model dominates M_0 on the images that appear almost indistinguishable from the clean image to human eyes.

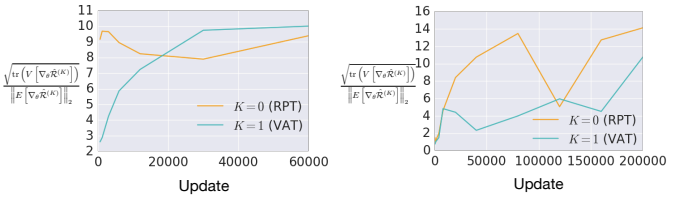

(a) MNIST (b) CIFAR-10

Fig. 7: Transition of the normalized SD norm of $\mathcal{R}^{(0)}$ and $\mathcal{R}^{(1)}$ during VAT training of NNs for supervised learning on MNIST and semi-supervised learning on CIFAR-10.

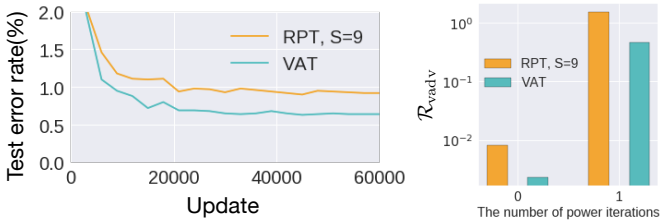

(a) Test error rate (b) $\mathcal{R}_{\mathrm{vadv}}$

Fig. 8: (a) Learning curves and (b) R_{vadv} of VAT implemented with $\alpha = 1$ and $S = 1$, and those of RPT implemented with optimal $\alpha(= 7)$ and $S = 9$. The remaining hyperparameter ϵ was set to 2.0 for both VAT and RPT.

REFERENCES

[1] Martın Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.

[2] Hirotugu Akaike. Information theory and an extension of the maximum likelihood principle. In *Selected Papers of Hirotugu Akaike*, pages 199–213. Springer, 1998.

[3] Philip Bachman, Ouais Alsharif, and Doina Precup. Learning with pseudo-ensembles. In *NIPS*, 2014.

[4] Chris M Bishop. Training with noise is equivalent to tikhonov regularization. *Neural computation*, 7(1):108–116, 1995.

[5] Ronan Collobert, Fabian Sinz, Jason Weston, and Léon Bottou. Large scale transductive svms. *Journal of Machine Learning Research*, 7(Aug):1687–1712, 2006.

[6] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. 2016.

[7] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In *AISTATS*, 2011.

[8] Gene H Golub and Henk A Van der Vorst. Eigenvalue computation in the 20th century. *Journal of Computational and Applied Mathematics*, 123(1):35–65, 2000.

[9] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. Deep learning. Book in preparation for MIT Press, 2016.

[10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014.

[11] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *ICLR*, 2015.

[12] Yves Grandvalet, Yoshua Bengio, et al. Semi-supervised learning by entropy minimization. In *NIPS*, 2004.

[13] Shixiang Gu and Luca Rigazio. Towards deep neural network architectures robust to adversarial examples. In *ICLR workshop*, 2015.

[14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. 2016.

[15] Gao Huang, Zhuang Liu, and Kilian Q Weinberger. Densely connected convolutional networks. *arXiv preprint arXiv:1608.06993*, 2016.

[16] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.

[17] Kevin Jarrett, Koray Kavukcuoglu, Marc'Aurelio Ranzato, and Yann LeCun. What is the best multi-stage architecture for object recognition? In *ICCV*, 2009.

[18] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.

[19] Diederik Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised learning with deep generative models. In *NIPS*, 2014.

[20] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. 2009.

[21] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *ICLR*, 2017.

[22] Chen-Yu Lee, Saining Xie, Patrick Gallagher, Zhengyou Zhang, and Zhuowen Tu. Deeply-supervised nets. In *AISTATS*, 2015.

[23] Min Lin, Qiang Chen, and Shuicheng Yan. Network in network. 2014.

[24] Lars Maaløe, Casper Kaae Sønderby, Søren Kaae Sønderby, and Ole Winther. Auxiliary deep generative models. In *ICML*, 2016.

[25] Andrew L Maas, Awni Y Hannun, and Andrew Y Ng. Rectifier nonlinearities improve neural network acoustic models. In *Proc. ICML*, volume 30, 2013.

[26] Shin-ichi Maeda. A bayesian encourages dropout. *arXiv preprint arXiv:1412.7003*, 2014.

[27] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, Ken Nakae, and Shin Ishii. Distributional smoothing with virtual adversarial training. In *ICLR*, 2016.

[28] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *ICML*, 2010.

[29] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011.

[30] Antti Rasmus, Mathias Berglund, Mikko Honkala, Harri Valpola, and Tapani Raiko. Semi-supervised learning with ladder networks. In *NIPS*, 2015.

[31] Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *NIPS*, 2016.

[32] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training GANs. 2016.

[33] Jost Tobias Springenberg. Unsupervised and semi-supervised learning with categorical generative adversarial networks. In *ICLR*, 2015.

[34] Jost Tobias Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin Riedmiller. Striving for simplicity: The all convolutional net. 2015.

[35] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *JMLR*, 15(1), 2014.

[36] Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. Highway networks. *arXiv preprint arXiv:1505.00387*, 2015.

[37] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *ICLR*, 2014.

[38] Theano Development Team. Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints*, abs/1605.02688, May 2016.

[39] Andrej Nikolaevich Tikhonov and Vasiliy Yakovlevich Arsenin. Solutions of ill-posed problems. 1977.

[40] Seiya Tokui, Kenta Oono, Shohei Hido, and Justin Clayton. Chainer: a next-generation open source framework for deep learning. In *Proceedings of workshop on machine learning systems (LearningSys) in the twenty-ninth annual conference on neural information processing systems (NIPS)*, 2015.

[41] Stefan Wager, Sida Wang, and Percy S Liang. Dropout training as adaptive regularization. In *NIPS*, 2013.

[42] Sumio Watanabe. Algebraic geometry and statistical learning theory. *Cambridge University Press*, 2009.

[43] Junbo Zhao, Michael Mathieu, Ross Goroshin, and Yann Lecun. Stacked what-where auto-encoders. 2016.

[44] Xiaojin Zhu and Zoubin Ghahramani. Learning from labeled and unlabeled data with label propagation. Technical report, Citeseer, 2002.

APPENDIX A
SUPERVISED CLASSIFICATION FOR THE MNIST DATASET

The MNIST dataset consists of 28×28 pixel images of handwritten digits and their corresponding labels. The input dimension is therefore $28 \times 28 = 784$, and each label is one of the numerals from 0 to 9. The following list summarizes the ranges for which we searched for the best hyperparameters of each regularization method:

- RPT: $\epsilon = [1.0, 50.0]$,
- adversarial training (with L_∞ norm constraint): $\epsilon = [0.05, 0.1]$,
- adversarial training (with L_2 norm constraint): $\epsilon = [0.05, 5.0]$, and
- VAT: $\epsilon = [0.05, 5.0]$.

All experiments were conducted with $\alpha = 1$ except when checking the effects of α in Section 4.2. Training was conducted using a mini-batch SGD based on ADAM [18]. We chose the mini-batch size of 100 and used the default values of [18] for the tunable parameters of ADAM. We trained the NNs with 60,000 parameter updates. For the base learning rate in validation, we selected the initial value of 0.002, and adopted the schedule of exponential decay with rate 0.9 per 600 updates. After the hyperparameter determination, we trained the NNs with with the same hyperparameters.

APPENDIX B
SUPERVISED CLASSIFICATION FOR CIFAR-10 DATASET

The CIFAR-10 dataset consists of $32 \times 32 \times 3$ pixel RGB images of objects categorized by their corresponding labels (cars, trucks, planes, animals, and humans). The number of training examples and test examples in the dataset are 50,000 and 10,000 respectively. We used 10,000 out of 50,000 training examples for validation and we applied ZCA whitening prior to the experiment. We also augmented the training dataset by applying a random 2×2 translation and random horizontal flip. We trained the Conv-Large model 6 over 300 epochs with batch size 100. For training, we used ADAM with essentially the same learning rate schedule as that used in [32]. In particular, we set the initial learning rate of ADAM to be 0.003 and linearly decayed the rate over the last half of training.

APPENDIX C
SEMI-SUPERVISED CLASSIFICATION FOR THE MNIST DATASET

For semi-supervised learning of MNIST, we used the same network as that used for supervised learning; however, we added zero-mean Gaussian random noise with 0.5 standard deviation to the hidden variables during training. This modification stabilized training on semi-supervised learning with VAT. We experimented with two sizes of labeled training samples, N_l = 100, 1000, and observed the effect of N_l on the test error. We used the validation set of fixed size 1,000, and used all the training samples, excluding the validation set and labeled training samples, to train the NNs; that is, when N_l = 100, the unlabeled training set N_{ul} had the size of $60,000 - 100 - 1,000 = 58,900$.

We searched for the best hyperparameter ϵ from $[0.05, 10.0]$. All experiments were conducted with $\alpha = 1$ and $K = 1$. For the optimization method, we again used ADAM-based mini-batch SGD with the same hyperparameter values as those in the supervised setting. We note that the likelihood term can be computed from labeled data only.

We used two separate mini-batches at each step: one mini-batch of size 64 from labeled samples to compute the likelihood term, and another mini-batch of size 256 from both labeled and unlabeled samples to compute the regularization term. We trained the NNs over 100,000 parameter updates, and started to decay the learning rate of ADAM linearly after we 50,000-th update.

APPENDIX D
SEMI-SUPERVISED CLASSIFICATION FOR THE SVHN AND CIFAR-10 DATASETS

The SVHN dataset consists of $32 \times 32 \times 3$ pixel RGB images of house numbers and their corresponding labels (0–9). The number of training samples and test samples within the dataset are 73,257 and 26,032, respectively. We reserved a 1,000 sample dataset for validation. From the remainder, we used a 1,000 sample dataset as a labeled dataset in semi-supervised training. Likewise in the supervised learning, we conducted ZCA preprocessing prior to the semi-supervised learning of CIFAR-10. We also augmented the training datasets with a random 2×2 translation. For CIFAR-10 exclusively, we also applied a random horizontal flip as well. For the labeled dataset, we used 4,000 samples randomly selected from the training dataset, from which we selected 1,000 samples for validation. We repeated the experiment three times with a different choice of labeled and unlabeled datasets on both SVHN and CIFAR-10.

For each benchmark dataset, we decided on the value of the hyperparameter ϵ based on the validation set. We also used a mini-batch of size 32 for the calculation of the negative log-likelihood term and used a mini-batch of size 128 for the calculation of $\mathcal{R}_{\mathrm{vadv}}$ in Eq. (8). We trained each model with 48,000 updates. This corresponds to 84 epochs for SVHN and 123 epochs for CIFAR-10. We used ADAM for training. We set the initial learning rate of ADAM to 0.001 and linearly decayed the rate over the last 16,000 updates. The performance of CNN-Small and CNN-Large that we reported in Section 4.1.2 is all based on training with data augmentation and the choices of ϵ that we described above.

TABLE 6: CNN models used in our experiments on CIFAR-10 and SVHN, based on [21], [32], [34]. All the convolutional layers and fully connected layers are followed by batch normalization [16] except the fully connected layer on CIFAR-10. The slopes of all lReLU [25] functions in the networks are set to 0.1.

Conv-Small on SVHN	Conv-Small on CIFAR-10	Conv-Large
32×32 RGB image		
3×3 conv. 64 lReLU	3×3 conv. 96 lReLU	3×3 conv. 128 lReLU
3×3 conv. 64 lReLU	3×3 conv. 96 lReLU	3×3 conv. 128 lReLU
3×3 conv. 64 lReLU	3×3 conv. 96 lReLU	3×3 conv. 128 lReLU
2×2 max-pool, stride 2 dropout, $p = 0.5$		
3×3 conv. 128 lReLU	3×3 conv. 192 lReLU	3×3 conv. 256 lReLU
3×3 conv. 128 lReLU	3×3 conv. 192 lReLU	3×3 conv. 256 lReLU
3×3 conv. 128 lReLU	3×3 conv. 192 lReLU	3×3 conv. 256 lReLU
2×2 max-pool, stride 2 dropout, $p = 0.5$		
3×3 conv. 128 lReLU	3×3 conv. 192 lReLU	3×3 conv. 512 lReLU
1×1 conv. 128 lReLU	1×1 conv. 192 lReLU	1×1 conv. 256 lReLU
1×1 conv. 128 lReLU	1×1 conv. 192 lReLU	1×1 conv. 128 lReLU
global average pool, 6×6 → 1×1		
dense 128 → 10	dense 192→ 10	dense 128→ 10
10-way softmax		

On SVHN, we tested the performance of the algorithm with and without data augmentation, and used the same setting that we used in the validation experiments for both Conv-Small and Conv-Large. For CIFAR-10, however, the models did not seem to converge with 48,000 updates; thus we reported the results with 200,000 updates.